

Checksums

Using DDT

Author: William Cooke (DRC/HPTg)
Phone: +1 - 609 - 987 - 5054
Fax: +1 - 609 - 987 - 5063
Email: William.Cooke@noaa.gov

National Oceanic and Atmospheric Administration
Geophysical Fluid Dynamics Laboratory
Princeton, NJ 08542
<http://www.gfdl.noaa.gov>



Outline

- Overview
- Checksums
- ddt



Outline

- Overview
- Checksums
- ddt



Overview

- Given an array of numbers, can we expect a simple summation to represent the true value of every element of the array?
- Is there a way to quickly tell when values change?
- Is there a tool to show where the code is failing?

Outline

- Overview
- Checksums
- ddt



What is a checksum?

[http://wordnetweb.princeton.edu/perl/webwn?
s=checksum](http://wordnetweb.princeton.edu/perl/webwn?s=checksum)

Noun

S: (n) checksum (a digit representing the sum of the digits in an instance of digital data; used to check whether errors have occurred in transmission or storage)

Precision of numbers

Given a model of number representation, how different in magnitude can numbers be before one of the numbers is considered irrelevant?

Debug session

- Open a debug session
 - `msub -X -q debug -I -l`
`partition=t1,size=NPES,`
`walltime=HH:MM:SS`
 - X allows X-windows to display back
 - I = interactive
 - l = following are submit commands
- Do this in a `freenx` window for X response times.

Precision exercise setup

- `msub -X -q debug -I -l partition=t1,size=32,walltime=02:00:00`
- `cd $CSCRATCH/$USER`
- `mkdir DDT`
- `cd DDT`
- `tar -xf ~William.Cooke/ddt_examples.tar`

Precision exercise

- cd precision
- Inspect the code precision.F90
 - Simple code that adds $1e-n$ to a scalar.
 - Prints out running total of the scalar
- Exercise
 - `./run_precision_r8.csh`
 - `./run_precision.csh`

Results of adding 10^{-n}

10^{-0}	1.00000000000000000000	10^{-10}	1.11111111110000000484
10^{-1}	1.10000000000000000888	10^{-11}	1.1111111111000000493
10^{-2}	1.11000000000000000977	10^{-12}	1.111111111100001382
10^{-3}	1.1109999999999999876	10^{-13}	1.11111111110000582
10^{-4}	1.1110999999999999766	10^{-14}	1.1111111111000502
10^{-5}	1.1111000000000000421	10^{-15}	1.11111111110001605
10^{-6}	1.1111109999999999598	10^{-16}	1.11111111110001605
10^{-7}	1.1111110000000000182	10^{-17}	1.11111111110001605
10^{-8}	1.1111111099999999574	10^{-18}	1.11111111110001605
10^{-9}	1.1111111100000000402	10^{-19}	1.11111111110001605

	Adding 1e-15		Adding 1e-16
1	1.00000000000000011	1	1.0000000000000000
2	1.00000000000000022	2	1.0000000000000000
3	1.00000000000000033	3	1.0000000000000000
4	1.00000000000000044	4	1.0000000000000000
5	1.00000000000000056	5	1.0000000000000000
6	1.00000000000000067	6	1.0000000000000000
7	1.00000000000000078	7	1.0000000000000000
8	1.00000000000000089	8	1.0000000000000000
9	1.00000000000000100	9	1.0000000000000000
10	1.00000000000000111	10	1.0000000000000000

Model of number format

http://en.wikipedia.org/wiki/Computer_number_format

IEEE 754-2008 standard defines a 64 bit floating point format with

11-bit exponent

52-bit significand

A sign bit

Real*8 (float) number =

$(\text{sign}) * (1 + \text{fractional significand}) * 2^{(\text{exponent} - 1023)}$

Model of number format

IEEE 754-2008 allows non zero number

- $\pm 1.797693134862231\text{E}+308$
- $\pm 4.940656458412465\text{E}-324$

Only good to 15 decimal places.

64 bits can be represented by 16 hexadecimal digits.

e.g. $1023(\text{base}10) = 11-1111-1111 \text{ (binary)} = 3\text{FF}$

$1.0 = \text{sign} * (1 + \text{significand} = 0) * 2^{(\text{exponent} = 1023 - 1023)}$

So $1.0 = 3\text{FF}0 \ 0000 \ 0000 \ 0000$

Model of number format

We still have the issue of trying to add 2 numbers together that are of vastly different magnitude.

Do this via the TRANSFER function.

- Use a long-integer (I^*8)
- Bitwise representation of the argument.
- Do an integer sum over the elements of an array.

Checksums in FMS models

Note that if 2 elements of an array are switched the checksum will be identical. However as climate models are chaotic, a small change due to a bug will cause answers to diverge.

These divergences can be spotted using the checksum capability of FMS.

&coupler_nml : do_chksum=.true.

This prints out the checksums of various types before and after a subroutine is called in the main loop.

Comparison of output will allow you to narrow down where the divergence is occurring.

NB : The output is verbose. Use only if absolutely necessary.



Outline

- Overview
- Checksums
- ddt



What is ddt?

- ddt is a graphical debugger.
- Similar to totalview.
- Compile code with '`-g -O0`'
- Do you really need a graphical debugger?
 - Maybe not. Examine the output of the program when run.

How to use?

Run in a directory that is visible to batch node

- `/lustre/fs` (Not `/lustre/lufs`!)

Load the ddt module

- `module load ddt`
- `ddt ./your_executable.x`
- `ddt -n $npes ./your_executable.x`
- You don't need to add `aprun`.

CSTARTMPI exercise

```
cd cstartmpi
```

```
./run_cstartmpi.csh
```

Runs program cstartmpi.exe

- without arguments (runs to completion)

- with arguments (fails)

- without arguments but a different pe count (fails with Segmentation fault)

Opening Window

You may
change
arguments.

Press run.

DDT - Run (on c2-login1)

Application: /lustre/fs/scratch/William.Cooke/20120628/cstartmpi/cstartmpi.exe foo bar blah Details ▲

Application: /lustre/fs/scratch/William.Cooke/20120628/cstartmpi/cstartmpi.exe Folder icon

Arguments: foo bar blah Dropdown arrow

Input File: Folder icon

Working Directory: Folder icon

☒ **MPI:** 4 processes, Cray XT/XE/XK (MPI/shmem) Details ▲

Number of processes: 4 Spinners

Implementation: Cray XT/XE/XK (MPI/shmem), no queue Change...

aprun arguments Dropdown arrow

☐ **OpenMP** Details ▼

☐ **CUDA** Details ▼

☐ **Memory Debugging** Details... ▼

Environment Variables: DMALLOC_OPTIONS debug=0x008 Details ▼

Plugins: none Details ▼

Run Cancel

Source Window

Allinea DDT v3.1-20638 (on c2-login1)

Session Control Search View Help

Current Group: All Focus on current: Group Process Thread Step Threads Together

Create Group

Project Files

Search (Ctrl+K)

Project Files

Source Tree

Header Files

Source Files

```
79
80 void (*s)(int);
81
82 MPI_Status status; /* Return status for receive */
83
84 t2 = malloc(sizeof(typeThree));
85
86 for(p=0;p<100;p++)
87     bigArray[p]=80000+p;
88
89 MPI_Init(&argc, &argv);
90 MPI_Comm_rank(MPI_COMM_WORLD, &my_rank);
91 MPI_Comm_size(MPI_COMM_WORLD, &p);
92
93 dynamicArray = malloc(sizeof(int)*100000);
94 sdin = malloc(sizeof(int) * p);
95
96 for(x=0;x<10000;x++)
97 {
98     dynamicArray[x] = x%10;
99 }
100
101 printf("my rank is %d\n", my_rank);
102
```

Locals Current Line(s) Current Stack

Current Line(s)

Variable Name	Value
argc	4
argv	0x7fffffffa08

Input/Output Breakpoints Watchpoints Stacks Tracepoints Tracepoint Output

Stacks

Processes	Function
4	main (cstartmpi.c:89)

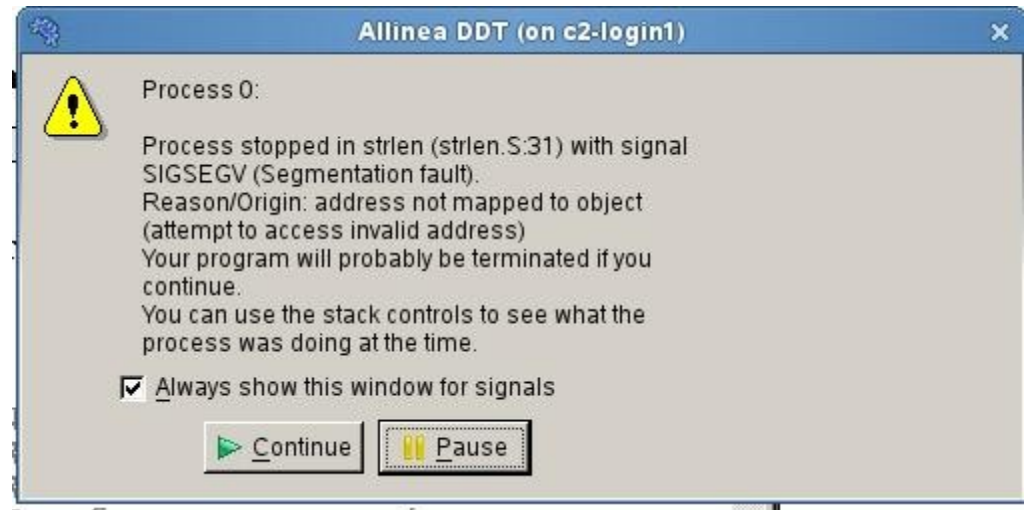
Evaluate

Expression	Value
------------	-------

Ready

SIGSEGV error message

There is an error

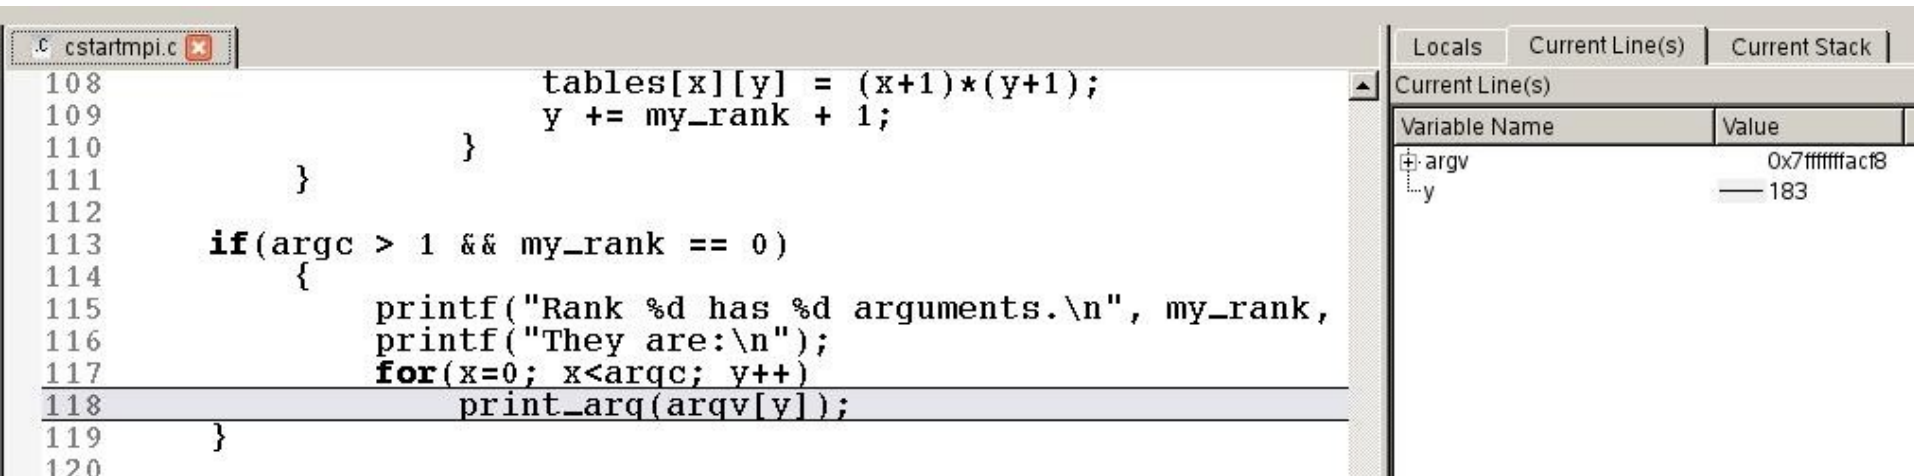


Press Pause

Look at the “Current Stack” window

Current Line Tab

Click on the "Current Line(s)" tab. What is the value of y?



The screenshot shows a debugger window with a C program named `cstartmpi.c`. The code is as follows:

```
108         tables[x][y] = (x+1)*(y+1);
109         y += my_rank + 1;
110     }
111 }
112
113 if(argc > 1 && my_rank == 0)
114 {
115     printf("Rank %d has %d arguments.\n", my_rank,
116     printf("They are:\n");
117     for(x=0; x<argc; y++)
118         print_arg(argv[y]);
119 }
120
```

The "Current Line(s)" tab is selected, showing the current state of the program. The variable `y` is highlighted, and its value is 183. The variable `argv` is also visible, with a value of `0x7fffffffa8`.

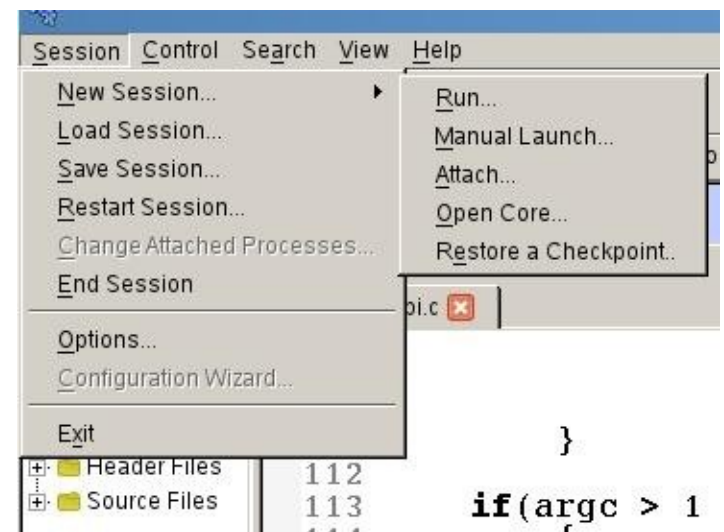
Variable Name	Value
argv	0x7fffffffa8
y	183

Is the value of y correct?

Is there anything wrong with this loop?

New Session

Correct the loop and run
./run_cstartmpi.csh again to
recompile the code.
Session-> New Session-> Run



Run the code to completion

Exercise 2

Session-> New Session-> Run
Change the core count to > 4

☒ MPI: 5 processes, Cray XT/XE/XK (MPI/shmem) Details ▲

Number of processes:

Implementation: Cray XT/XE/XK (MPI/shmem), no queue Change...

aprun arguments

Find out what is
the problem.

```
    y = 0;
    while(y != 12)
    {
        tables[x][y] = (x+1)*(y+1);
        y += my_rank + 1;
    }

    > 1 && my_rank == 0
    printf("Rank %d\n", my_rank);
    printf("Thread %d\n", thread_id);
    for(x=0; x<12; x++)
    for(y=0; y<12; y++)
    printf("x=%d y=%d\n", x, y);
```

Allinea DDT (on c2-login1)

Process 4:

Memory error detected in main (cstartmpi.c:108).
Process attempted to dereference a null pointer or execute an SSE instruction with an incorrectly aligned memory address (the latter may sometimes occur spuriously if guard pages are enabled)
Tip: Use the stack list and the local variables to explore your program's current state and identify the source of the error.

Continue Pause

trisol exercise

```
cd trisol
```

```
./run_trisol.csh
```

Runs program trisol.exe

Memory error

Walkthrough of how to turn on memory debugging.

trisol exercise

```
> aprun -n 4 ./trisol.exe
```

```
*** Solution correct
```

```
|x| / (sqrt(n)*epsilon*(|A|*|x| + |b|) = 3.7352-215
```

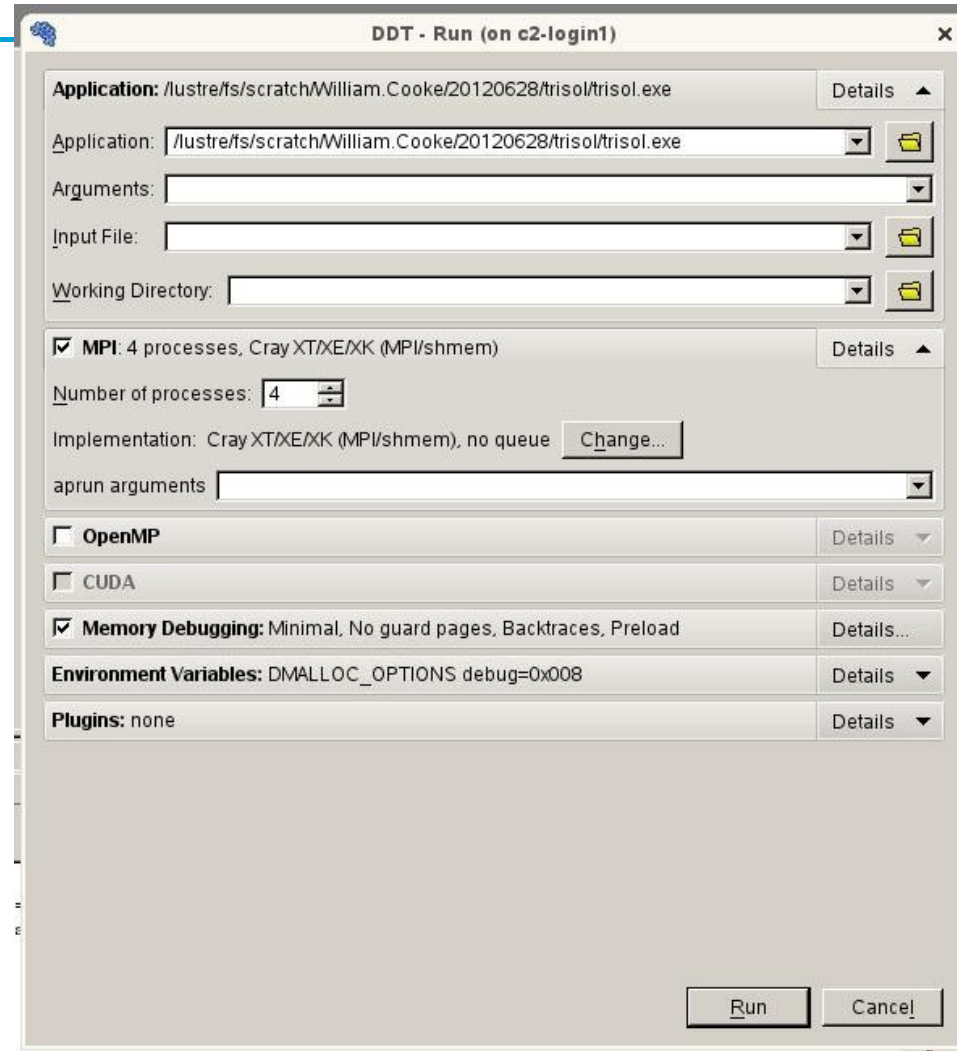
```
> ddt -n 4 ./trisol.exe
```

Runs to completion.

So everything is good?? No. There is a memory error in here.

Memory Debugging

Restart the session.
Check the Memory
Debugging box.
Click on Details.



DDT - Run (on c2-login1)

Application: /lustre/fs/scratch/William.Cooke/20120628/trisol/trisol.exe Details ▲

Application: /lustre/fs/scratch/William.Cooke/20120628/trisol/trisol.exe

Arguments:

Input File:

Working Directory:

☒ **MPI:** 4 processes, Cray XT/XE/XK (MPI/shmem) Details ▲

Number of processes: 4

Implementation: Cray XT/XE/XK (MPI/shmem), no queue Change...

aprun arguments:

☐ **OpenMP** Details ▼

☐ **CUDA** Details ▼

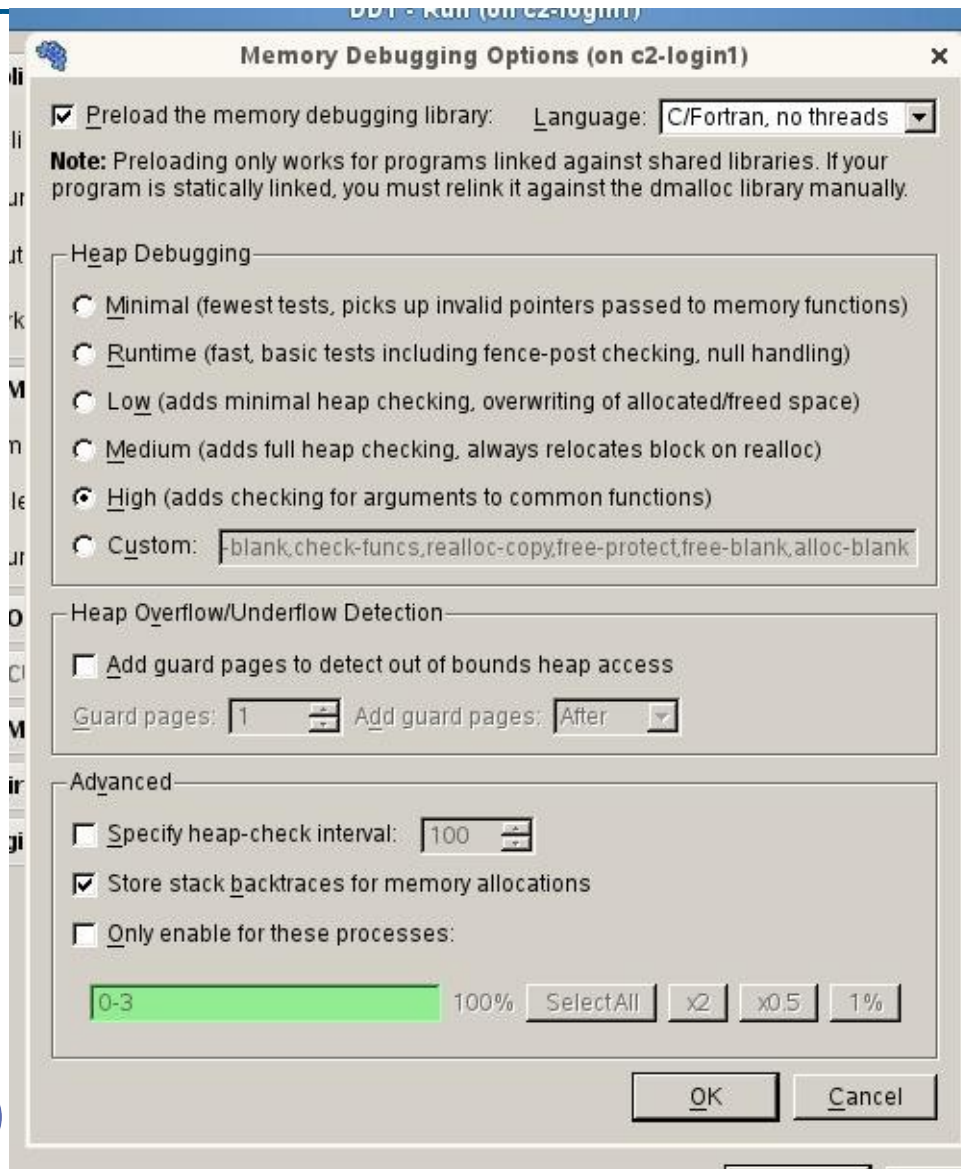
☒ **Memory Debugging:** Minimal, No guard pages, Backtraces, Preload Details...

Environment Variables: DMALLOC_OPTIONS debug=0x008 Details ▼

Plugins: none Details ▼

Run Cancel

Memory Debugging Options



Check Preload the memory.

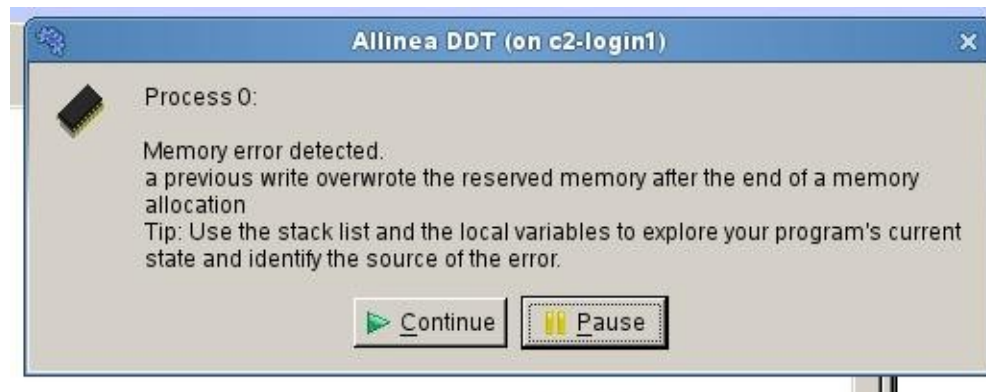
In the Heap debugging options check the High box.

Leave the Heap Overflow unchecked (for now).

Memory Error.

Run the program.

DDT reports an array overflow.

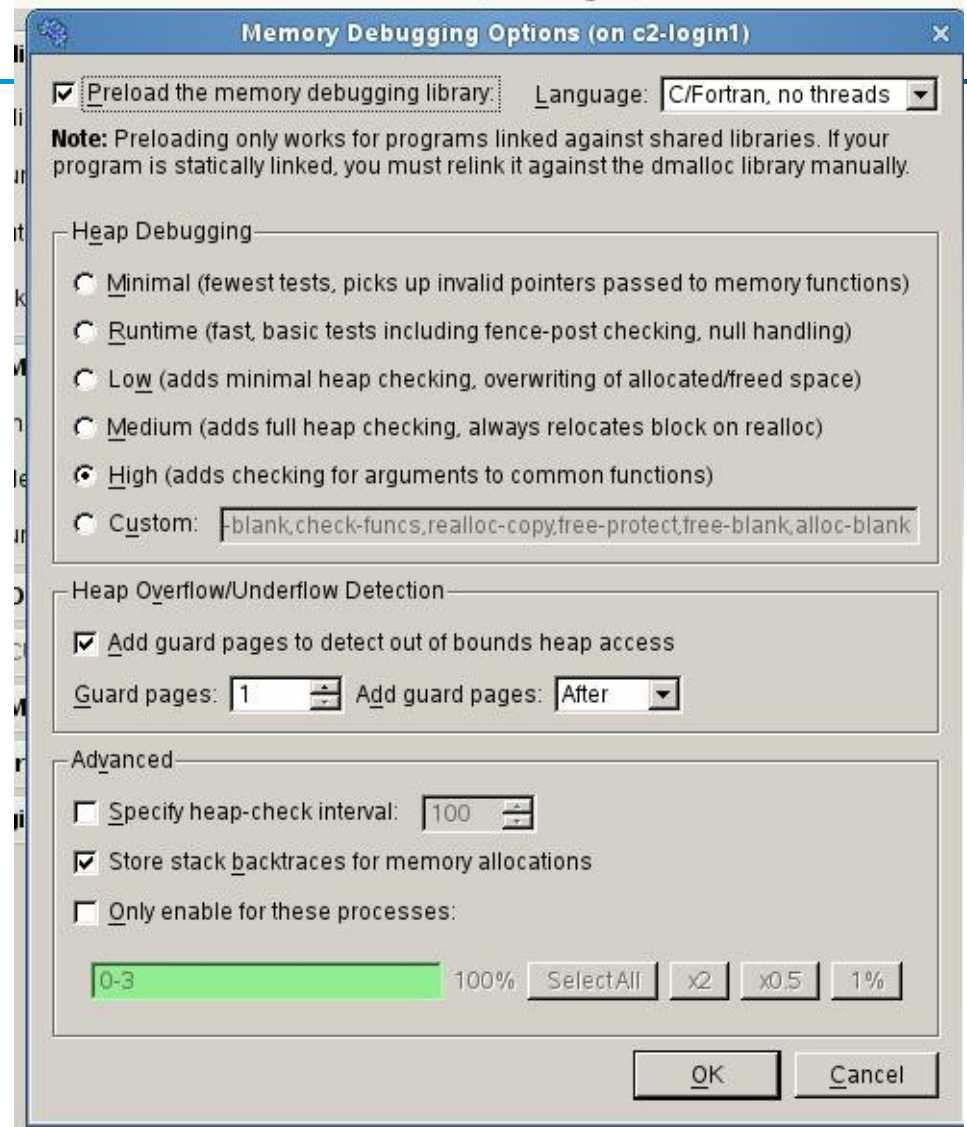


Where is the overflow?

We'll use guard pages to find out.

Guard Pages

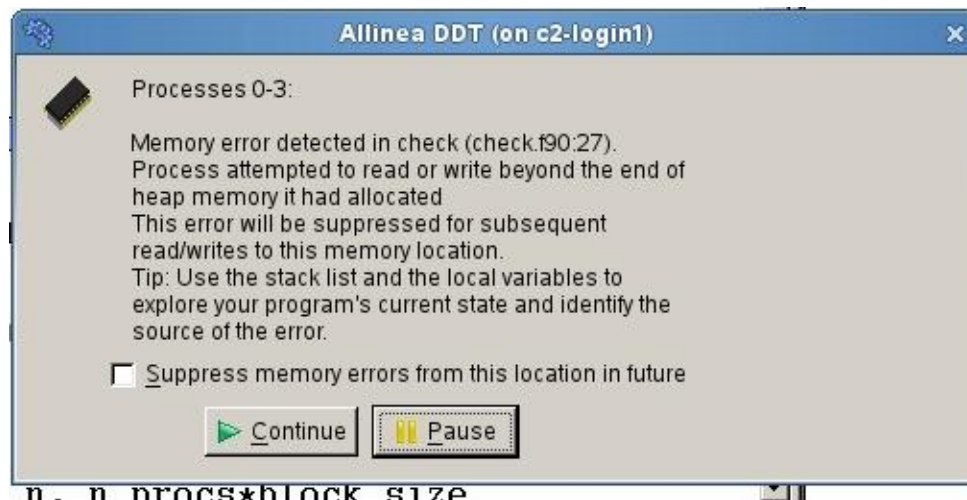
Restart the session.
Click on Memory
Debugging->Details.
Check the "Add
guard pages" box
under "Heap
Overflow/Underflow
Detection"
Leave at "1" and
"After"



Memory Error.

Run the program again.

DDT reports a memory error in check.f90



Click Pause

Look at the size of res in the "Current Lines" tab.

Size of res.

0 1 2 3

```
sol.f90 x f check.f90 x
n = SIZE(a,1)
ALLOCATE (res(n),temp(n))

IF (me.LE.0) THEN
    temp = b
ELSE
    temp = ZERO
END IF

DO k = 0, block_size
    res(k+1) = k+1
END DO
```

Locals	Current Line(s)	Current Stack
Current Line(s)		
Variable Name	Value	
k	4095	
res		
Type: REAL*8 (4095)		

res is size 4095

k = 4095 so we are trying to write beyond
the end of res (res(k+1))

Exercise 3.

```
Fix the code by inserting  
DO k = 1, block_size  
    res(k) = k  
END DO
```

And run `./run_trisol.csh` to recompile
Then
`ddt -n 4 ./trisol.exe`

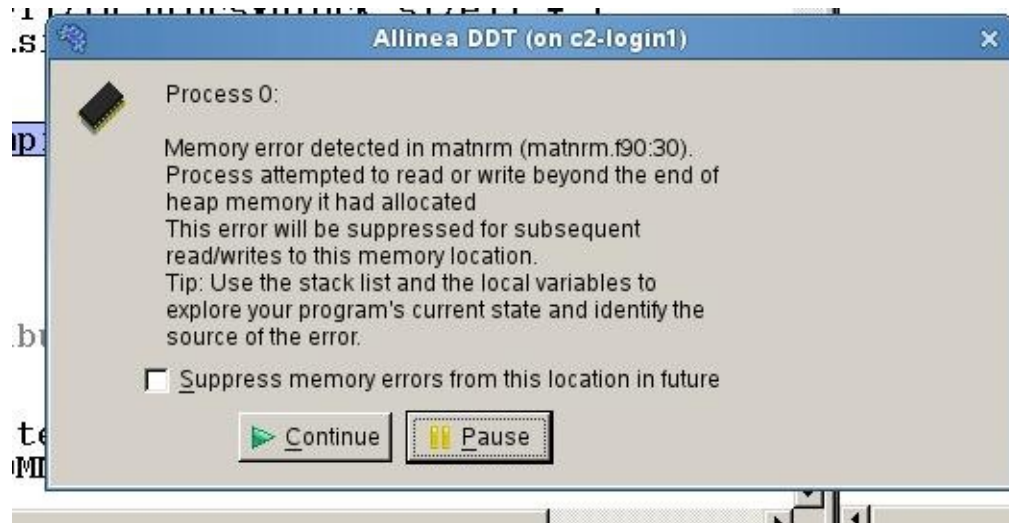
Exercise 3 continued.

You should get a
Memory Error in
matnrm.

Pause the job and
examine values of i,
k, j1, j2 on different
processors.

You'll need to pause all cores.

See if you can figure out what is wrong.



Exercise 3 Hints

What is the `block_size` on each core?

Does it seem correct?

Where is it defined?

How do you define an array in Fortran?

Questions?

